

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Exclusive Encryption

Inventor(s):

John R. Douceur

Atul Adya

Josh D. Benaloh

Gideon A. Yuval

ATTORNEY'S DOCKET NO. MS1-712US

1 **TECHNICAL FIELD**

2 This invention relates to cryptography, and more particularly to
3 cryptosystems that enable enforcement of syntactical and/or semantic
4 requirements on encrypted data.
5

6 **BACKGROUND OF THE INVENTION**

7 As computer technology has advanced, so too has the need for computer
8 security. One significant area of concern with computer security is the security of
9 data stored by a user of the computer. Such security can include, for example,
10 security against another user reading the data, another user modifying the data
11 (either maliciously or unintentionally), etc. One manner in which such security is
12 currently provided is through encryption. Using encryption, only those users who
13 are supposed to have access to the data (e.g., those users with the decryption key)
14 are able to decrypt and read (and/or modify) the data.

15 However, situations arise in which encryption creates additional problems.
16 One such situation is in the case where the data is to be verified as being in
17 accordance with a certain set of rules (e.g., a naming syntax), while at the same
18 time maintaining the security/privacy of the data. Current systems typically
19 cannot perform such verifications without reading the underlying data.

20 For example, in some situations where computers are networked together it
21 would be desirable to have a distributed file system in which different files could
22 be stored on different machines in the network. However, in order to maintain
23 security of the data in the files, it would be beneficial to be able to render both the
24 data and the filename itself unreadable at the computer on which the file is stored,
25 thereby preventing another unauthorized user from gathering information about

1 the data by either reading the data file or the filename. However, maintaining a
2 directory of such encrypted/protected filenames can lead to difficulties, especially
3 due to the fact that the user of the computer on which the data file or filename is
4 stored may not be authorized to read the file. Care must be taken to ensure that the
5 filename syntax is not violated and that duplicate filenames do not exist. Current
6 systems do not have the ability to maintain such a directory of non-duplicate
7 filenames that do not violate the naming syntax while at the same time can be
8 verified by a component that is not authorized to read the filenames.

9 The invention described below addresses these disadvantages, providing
10 methods and systems for encryption that excludes syntactically illegal plaintext
11 from being encrypted and that enables a party without access to encryption keys to
12 exclude more than one item of ciphertext that decrypts to the same plaintext. For
13 this reason, the invention is referred to as "exclusive encryption."

14 15 **SUMMARY OF THE INVENTION**

16 Methods and systems for exclusive encryption are described herein.

17 According to one aspect, an exclusive encryption system is established
18 using multiple computing devices. The exclusive encryption system allows for the
19 exclusion of certain plaintext (e.g., by one of the computing devices) and
20 ciphertext (e.g., by another of the computing devices) while at the same time
21 maintaining the privacy created by the encryption (e.g., so the other computing
22 device cannot see the plaintext).

23 According to another aspect, an exclusive encryption system is established
24 as part of a serverless distributed file system. The file system is distributed among
25 multiple computing devices, and each directory entry (e.g., a file name or folder

1 name) is encrypted by the device creating (or otherwise using) the entry. The
2 encryption process used ensures that the directory entry is syntactically legal and
3 cannot be read by the device on which the entry is stored. Additionally, the device
4 on which the entry is stored is able to verify the legality of the directory entry, and
5 ensures that duplicate entries are not stored in the same directory.

6 7 **BRIEF DESCRIPTION OF THE DRAWINGS**

8 The present invention is illustrated by way of example and not limitation in
9 the figures of the accompanying drawings. The same numbers are used
10 throughout the figures to reference like components and/or features.

11 Fig. 1 is a block diagram illustrating an exemplary system in which certain
12 embodiments of the invention can be implemented.

13 Fig. 2 illustrates an exemplary network environment in which certain
14 embodiments of the invention may be practiced.

15 Fig. 3 illustrates an exemplary computing device such as may be used in
16 accordance with certain embodiments of the invention.

17 Fig. 4 is a flowchart illustrating an exemplary process for exclusively
18 encrypting a directory entry prior to storage in accordance with certain
19 embodiments of the invention.

20 Fig. 5 is a flowchart illustrating an exemplary process for decrypting a
21 received directory entry in accordance with certain embodiments of the invention.

22 Fig. 6 is a flowchart illustrating an exemplary process for mapping a
23 plaintext name to a mapped name in accordance with certain embodiments of the
24 invention.

Fig. 7 is a flowchart illustrating an exemplary process for decasifying a mapped name in accordance with certain embodiments of the invention.

Fig. 8 is a flowchart illustrating an exemplary process for encoding a decasified name in accordance with certain embodiments of the invention.

Fig. 9 is a flowchart illustrating an exemplary process for verifying a received directory entry in accordance with certain embodiments of the invention.

Fig. 10 is a flowchart illustrating an exemplary process for decoding an encoded name in accordance with certain embodiments of the invention.

Fig. 11 is a flowchart illustrating an exemplary recasifying procedure in accordance with certain embodiments of the invention.

Fig. 12 is a flowchart illustrating an exemplary process for demapping a name in accordance with certain embodiments of the invention

Fig. 13 shows a general example of a computer that can be used in accordance with the invention

DETAILED DESCRIPTION

Cryptographic methods and systems including exclusive encryption are described herein. Exclusive encryption is used herein to refer to encryption that allows for the exclusion of certain plaintext and ciphertext while at the same time maintaining the privacy created by the encryption. In certain embodiments, these exclusions include the exclusion of syntactically illegal plaintext (e.g., exclusion from being encrypted) and the exclusion of multiple ciphertexts that decrypt to the same plaintext.

The discussions herein assume a basic understanding of cryptography by the reader. For a basic introduction of cryptography, the reader is directed to a

1 text written by Bruce Schneier and entitled "Applied Cryptography: Protocols,
2 Algorithms, and Source Code in C," published by John Wiley & Sons with
3 copyright 1994 (or second edition with copyright 1996).

4 Fig. 1 is a block diagram illustrating an exemplary system in which certain
5 embodiments of the invention can be implemented. A computing device 102 is
6 illustrated coupled to a computing device 104. Computing device 102
7 communicates ciphertext 106 to computing device 104. Ciphertext 106 can be
8 communicated to computing device 104 for any of a wide variety of uses, such as
9 storage of the ciphertext, verification of the ciphertext, forwarding of the
10 ciphertext, etc. Ciphertext 106 is encrypted using an exclusive encryption
11 cryptosystem that allows computing device 104 to exclude certain types of data
12 (e.g., duplicates) from being processed (e.g., stored, forwarded, etc.) by device 104
13 without actually determining the underlying plaintext.

14 Another computing device 108 coupled to computing device 104 may also
15 be able to access, and obtain the ciphertext 106 from, computing device 104. If
16 computing device 108 is not authorized to access the plaintext corresponding to
17 ciphertext 106, then computing device 108 would not be able to decrypt ciphertext
18 106 (e.g., device 108 would not have the appropriate cryptographic key(s)) and the
19 corresponding plaintext would thus be secure from being viewed by device 108.
20 However, if computing device 108 is authorized to access the plaintext
21 corresponding to ciphertext 106 (e.g., computing device has the appropriate key(s)
22 to decrypt ciphertext 106), then device 108 will be able to obtain the plaintext by
23 decrypting ciphertext 106. Additionally, computing device 108 is assured that the
24 decrypted plaintext will be syntactically legal, regardless of what encrypted bits it
25 receives and regardless of what key(s) it uses.

09764962-01701

1 The methods and systems of the exclusive encryption cryptosystem are
2 discussed herein primarily with reference to a directory service that uses exclusive
3 encryption. However, the invention is not limited to use with directory services
4 and can be used in any of a wide variety of other situations. One example of such
5 a situation is a secure courier service in which a courier guarantees to the message
6 recipient that it will only deliver syntactically valid messages, and it will never
7 deliver a duplicate message, even though the courier is not permitted access to the
8 unencrypted contents of the messages it carries. Another example is a secret-
9 ballot election, in which one wishes to transmit a vote in encrypted form while
10 demonstrating that the vote is of a valid form. Yet another example is "secret
11 sharing", in which portions of a key are distributed; in such situations, one may
12 want to encrypt the "shares" of the key for transmission while demonstrating that
13 the shares are valid.

14 Fig. 2 illustrates an exemplary network environment 118 in which certain
15 embodiments of the invention may be practiced. Four client computing devices
16 120, 122, 124, and 126 are illustrated coupled together via a data communications
17 network 128. Although four computing devices 120 – 126 are illustrated, different
18 numbers (either greater or fewer than four) may be included in network
19 environment 118.

20 Network 128 represents any of a wide variety of data communications
21 networks. Network 128 can include public portions (e.g., the Internet) as well as
22 private portions (e.g., an internal corporate Local Area Network (LAN)), as well
23 as combinations of public and private portions. Network 128 can be implemented
24 using any one or more of a wide variety of conventional communications media
25 including both wired and wireless media. Any of a wide variety of

communications protocols can be used to communicate data via network 128, including both public and proprietary protocols. Examples of such protocols include TCP/IP, IPX/SPX, NetBEUI, etc.

Computing devices 120 – 126 represent any of a wide range of computing devices, and each device 120 – 126 may be the same or different. By way of example, devices 120 – 126 may be desktop computers, laptop computers, handheld or pocket computers, personal digital assistants (PDAs), cellular phones, Internet appliances, consumer electronics devices, gaming consoles, etc.

Two or more of devices 120 – 126 operate to implement a serverless distributed file system 130. The actual devices included in the serverless distributed file system 130 can change over time, allowing new devices to be added to the system 130 and other devices to be removed from the system 130. Each device 120 – 126 that is part of the distributed file system 130 has different portions of its mass storage device(s) (e.g., hard disk drive) allocated for use as either local storage or distributed storage. The local storage is used for data that the user desires to store on his or her local machine and not in the distributed file system structure. The distributed storage portion is used for data that the user of the device (or another device) desires to store within the distributed file system structure. In the illustrated example of Fig. 2, each of devices 120 – 124 includes one or more mass storage devices that each includes a distributed portion and a local portion. This separation into multiple portions may occur on a per storage device basis (e.g., one hard drive is designated for use in the distributed system 130 while another is designated solely for local use), and/or within a single storage device (e.g., part of one hard drive may be designated for use in the distributed system 130 while another part is designated for local use). Another computing

device, device 126, is coupled to network 128 but is not part of the distributed file system 130 and thus does not have any of its mass storage device(s) allocated for use by the distributed system. Device 126, which does not have storage space allocated for use as distributed storage, may be permitted to store files in the distributed portions of one or more of devices 120 – 124, or alternatively may be prohibited from doing so.

The distributed file system 130 operates to store one or more copies of files on different computing devices 120 – 124. When a new file is created by the user of a computer 120 – 126, he or she has the option of storing the file on the local portion of his or her computing device, or alternatively in the distributed file system 130. If the file is stored in the distributed file system 130, then the file will be stored in the distributed system portion of the mass storage device(s) of one or more of devices 120 – 124. The user creating the file typically has no ability to control which device 120 – 124 the file is stored on, nor any knowledge of which device 120 – 124 the file is stored on. Additionally, duplicate copies of the file will typically be saved, allowing the user to subsequently retrieve the file even if one of the computing devices 120 – 124 on which the file is saved is unavailable (e.g., is powered-down, is malfunctioning, etc.).

The distributed file system 130 is implemented by one or more components on each of the devices 120 – 124, thereby obviating the need for any centralized server to coordinate the file system. These components operate to determine where particular files are stored, how many copies of the files are created for storage on different devices, etc. Exactly which device will store which files depends on numerous factors, including the number of devices in the distributed file system 130, the storage space allocated to the file system from each of the

1 devices, a number of copies of the file are to be saved, the number of files already
2 stored on the devices, etc. Thus, the distributed file system 130 allows the user to
3 create and access files (as well as folders or directories) without any knowledge of
4 exactly which other computing device(s) the file is being stored on.

5 The files stored by the file system 130 are distributed among the various
6 devices 120 – 124 and stored in encrypted form. When a new file is created, the
7 device on which the file is being created encrypts the file prior to communicating
8 the file to other device(s) for storage. The directory entry (e.g., the file name) for
9 a new file is also communicated to the other device(s) for storage. Additionally, if
10 a new folder or directory is created, the directory entry (e.g., folder name or
11 directory name) is also communicated to the other device(s) for storage. As used
12 herein, a directory entry refers to any entry that can be added to a file system
13 directory, including both file names and directory (or folder) names.

14 The distributed file system 130 is designed to prevent unauthorized users
15 from reading data stored on one of the devices 120 – 124. Thus, a file created by
16 device 122 and stored on device 124 is not readable by the user of device 124
17 (unless he or she is authorized to do so). In order to implement such security, the
18 contents of files as well as all directory entries are encrypted, and only authorized
19 users are given the decryption key. Thus, although device 124 may store a file
20 created by device 122, if a user of device 124 is not an authorized user of the file
21 then the user of device 124 cannot decrypt (and thus cannot read) either the
22 contents of the file or its directory name (e.g., filename).

23 The exclusive encryption described herein allows the directory entry to be
24 stored in an encrypted form, thereby preventing unauthorized users from
25 improperly gaining any information based on the filename. Additionally, the

1 exclusive encryption assures that various properties are maintained. In the
2 illustrated example, the following three properties are maintained:

- 3 • no two encrypted entries in a directory will decrypt to the same
4 name
- 5 • all encrypted entries in the directory decrypt to syntactically legal
6 names
- 7 • the service that maintains the directory does not have access to the
8 plaintext names of the entries

9 Thus, the system is able to ensure both that no two entries in a directory are
10 encryptions of the same name and that all entries in a directory are encryptions of
11 syntactically legal names, while at the same time ensuring that the device
12 maintaining the directory does not have access to the plaintext names of the
13 entries.

14 Fig. 3 illustrates an exemplary computing device 150 such as may be used
15 in accordance with certain embodiments of the invention. Computing device 150
16 can be, for example, a device 120 – 124 of Fig. 2. Computing device 150 includes
17 a server component 152, a client component 154, a memory 156, and a mass
18 storage device 158. Computing device 150 also typically includes additional
19 components (e.g., a processor), however these additional components have not
20 been shown in Fig. 3 so as not to clutter the drawings. Memory 156 can be any of
21 a wide variety of conventional volatile and/or nonvolatile memories, such as
22 RAM, ROM, Flash memory, etc. Mass storage device 158 can be any of a wide
23 variety of conventional nonvolatile storage devices, such as a magnetic disk,
24 optical disk, Flash memory, etc.

1 Computing device 150 is intended to be used in a serverless distributed file
2 system, and as such includes both server component 152 and client component
3 154. Server component 152 handles requests when device 150 is responding to a
4 request involving a file or directory entry stored (or to be stored) in storage device
5 158 (although files and corresponding directory entries need not be stored on the
6 same device 150), while client component 154 handles the issuance of requests by
7 device 150 for files stored (or to be stored) in the distributed file system. Client
8 component 154 and server component 152 operate independent of one another.
9 Thus, situations can arise where the serverless distributed file system causes files
10 being stored by client component 154 to be stored in mass storage device 158 by
11 server component 152.

12 Client component 154 includes a storage and retrieval control module 160,
13 a mapping module 162, a decasifying module 164, an encoding module 166, an
14 encryption module 168, a decryption module 170, a decoding module 172, a
15 recasifying module 174, and a demapping module 176. Control module 160 (in
16 combination with interface 200 discussed below) manages accesses to the
17 serverless distributed file system for the creation, storage, retrieval, and any other
18 modification of files and directories on behalf of computing device 150. Due to
19 the distributed nature of the file system, directory information such as file names
20 and directory names (plaintext names) are modified by computing device 150
21 prior to being stored in the file system so that the modified names themselves
22 reveal very little information about the original plaintext name, while at the same
23 time allowing the device at which the names will be stored to verify that they are
24 syntactically legal names and are not duplicates. The modified names do leak
25 information about the approximate length of the original plaintext name (although

1 this leak can be avoided in an alternate embodiment discussed below). The
2 modules 162 – 168 are used to modify plaintext names prior to communicating
3 them to the serverless distributed file system, while modules 170 – 176 are used to
4 reconstruct the plaintext name when a modified name is received from the
5 serverless distributed file system. The use of these modules and their operation
6 will be discussed with additional reference to Figs. 4 – 12.

7 Fig. 4 is a flowchart illustrating an exemplary process for exclusively
8 encrypting a directory entry prior to storage in accordance with certain
9 embodiments of the invention. Initially, control module 160 receives a plaintext
10 name that is the directory entry to be encrypted (act 180). The plaintext name is
11 received from another component internal to device 150 (not shown), such as an
12 application, operating system, etc. The received plaintext name is made available
13 to mapper 162 which maps the plaintext name to a new mapped name (act 182).
14 The mapped name is then made available to decasifier 164 which decasifies the
15 mapped name into a decasified name (e.g., a name where the case of the characters
16 has no importance) and corresponding case information (act 184).

17 The decasified name is then made available to encoder 166 which encodes
18 the decasified name (act 186). The encoded name is then made available to
19 encryptor 168 which encrypts both the encoded name from act 186 and the case
20 information from act 184 (act 188). Control module 160 then communicates the
21 encrypted name and encrypted case information to another device for verification
22 and storage (act 190). This encrypted name and case information are also referred
23 to as ciphertexts.

24 The process of Fig. 4 is based on the directory entries being case-
25 insensitive. That is, the serverless distributed file system does not distinguish

1 between upper-case and lower-case letters. For example, the entries "MEMO" and
2 "memo" would be the same entries as viewed by the serverless distributed file
3 system. Alternatively, the serverless distributed file system may be case-sensitive.
4 If the system is case-sensitive, then no decasifying is done, and neither decasifier
5 164 nor decasifying act 184 are needed (and act 188 would not encrypt case
6 information).

7 After the directory entry is modified per the process of Fig. 4, the directory
8 entry is communicated to another device via a distributed system interface 200.
9 Distributed system interface 200 determines, optionally in collaboration with
10 corresponding interfaces on other devices in the serverless distributed file system,
11 where the directory entry and corresponding file (if any) are to be stored.
12 Interface 200 selects one or more devices, which can include computing device
13 150, and communicates the directory entry to the selected device. The selection
14 process can be accomplished in a variety of different manners, such as based on
15 the storage space available in the various devices that are part of the serverless
16 distributed file system, an anticipated likelihood that one or more of the various
17 devices will be unavailable at any given moment, etc. Interface 200 also
18 maintains a client storage record 202 that identifies the location of each directory
19 entry (and corresponding file, if any) that is stored by client component 154.
20 Thus, when client component 154 requests retrieval of a particular directory entry
21 interface 200 knows where to locate the requested entry.

22 Fig. 5 is a flowchart illustrating an exemplary process for decrypting a
23 received directory entry in accordance with certain embodiments of the invention.
24 Initially, control module 160 receives an encrypted name and encrypted case
25 information (ciphertexts) from another device via interface 200 (act 210). The

1 encrypted name and encrypted case information are made available to decryptor
2 170 which decrypts both the name and the case information (act 212). The
3 decrypted name is then made available to decoder 172 which decodes the
4 decrypted name (act 214).

5 The decoded name and decrypted case information are then made available
6 to recasifier 174 which uses the case information to recasify the decoded name
7 (act 216). The recasified name is then made available to demapper 176 which
8 demaps the recasified name into the plaintext name (act 218).

9 The process of Fig. 5 is based on the directory entries being case-
10 insensitive. Alternatively, if the serverless distributed file system is case-sensitive,
11 then no recasifying is done, and neither recasifier 174 nor recasifying act 216 are
12 needed (and no decryption of case information is performed in act 212).

13 An exemplary implementation of the modules 162 – 176 will now be
14 discussed.

15 Mapping module (mapper) 162 maps a received plaintext name into a
16 mapped name. Fig. 6 is a flowchart illustrating an exemplary process for mapping
17 the plaintext name to the mapped name in accordance with certain embodiments of
18 the invention. Mapping module 162 initially receives the input plaintext name (act
19 240). Mapping module 162 may receive the actual name, or alternatively an
20 indication of where the name can be found (e.g., a pointer to the name).

21 Mapping module 162 then checks whether the input plaintext name is an
22 illegal name (act 242). Mapping module 162 includes a record of illegal names,
23 which may be pre-programmed or defined in module 162 or alternatively provided
24 to module 162 by some other component (e.g., module 160). In one
25 implementation, the following strings (regardless of case) are illegal names:

1 CONIN\$, CONOUT\$, CON, PRN, AUX, NUL, LPT n , COM n (where n
2 represents any digit 0 through 9). Alternatively, the illegal names may include
3 more or fewer (and/or different) strings. Additionally, the set of illegal names
4 may be the empty set, which is equivalent to there being no illegal names. If there
5 are no illegal names, then neither mapping module 162 (nor mapping act 242)
6 need be included, and demapping module 176 and any demapping acts (discussed
7 in more detail below) also need not be included.

8 If the plaintext name is equal to any one of these illegal names then an
9 indication that no valid mapping exists is given to module 160 (act 244) and the
10 modification process of Fig. 4 fails. However, if the input plaintext name is not an
11 illegal name, then module 162 checks whether the input plaintext name is an
12 illegal name followed by one or more underscore characters, such as the name
13 "CON_" or "PRN__", (act 246). If the input name is not an illegal name followed
14 by one or more underscore characters, then the input plaintext name is used as the
15 mapped name (act 248). In other words, the input name maps to itself. However,
16 if the input name is an illegal name followed by one or more underscore
17 characters, then the mapped name is the input name with one trailing underscore
18 removed (act 250). For example, the input name "CON_" would map to "CON",
19 and the input name "PRN__" (having two underscores) would map to "PRN_"
20 (having one underscore).

21 Alternatively, other characters may be used by mapping module 162 other
22 than an underscore. Any character (e.g., letter, number, other symbol, etc.) can be
23 used in place of the underscore, so long as adding a string of zero or more such
24 characters to an illegal name does not result in another illegal name.

1 Once mapped, the mapped name is decasified by decasifying module
2 (decasifier) 164. Fig. 7 is a flowchart illustrating an exemplary process for
3 decasifying the mapped name in accordance with certain embodiments of the
4 invention. Decasifying module 164 initially receives the input mapped name (act
5 270). Decasifying module 164 may receive the actual name, or alternatively an
6 indication of where the name can be found (e.g., a pointer to the name), from
7 module 160 or alternatively directly from mapping module 162.

8 Module 164 then initializes two strings, referred to as the *L* and the *I*
9 strings, to null (act 272). At the end of the decasifying process, the *L* string will be
10 the decasified version of the input name received in act 270, and the *I* string will
11 be a set of bits that identify the appropriate case for each of the characters in the *L*
12 string. Module 164 operates to decasify names based on the Unicode standard
13 including current and future versions of the standard. Additionally information
14 regarding the Unicode standard is available from The Unicode Consortium of
15 Mountain View, California. The Unicode standard uses sixteen bits to identify
16 each character, but it is an evolving standard, so not all 65,536 combinations are
17 defined. However, the conventional ASCII character set (which includes the
18 uppercase and lowercase English alphabet, numbers, and many symbols
19 commonly used in English) is defined using the last seven bits of the sixteen
20 Unicode bits (with the first nine bits being all zeroes). Thus, characters with
21 Unicode values between 0 and 127 are equivalent to characters with identical
22 ASCII values. Furthermore, characters with Unicode values between 128 and 255
23 have all been defined, and they include characters that, together with the first 128
24 Unicode characters, are sufficient to display text in most Latinate languages.
25 These 256 characters are sometimes referred to as the “extended ASCII” character

1 set. Due to the evolving nature of the Unicode standard, module 164 is
2 implemented so as to decasify only those characters with Unicode values less than
3 or equal to 255 (that is, the extended ASCII character set). Module 164 does not
4 attempt to decasify any character in the input mapped name having a Unicode
5 value greater than 255. Alternatively, additional characters may be decasified,
6 such as any characters which are defined in the Unicode standard.

7 To decasify the received name, module 164 selects the first character from
8 the input name (act 274). Module 164 then checks whether the leading eight bits
9 of the selected character are zero (act 276). If the leading eight bits are not all zero
10 (that is, the character has a Unicode value greater than 255) then module 164
11 appends the selected character to the *L* string and a zero bit to the *I* string (act
12 278). Module 164 then checks whether there are any additional characters in the
13 input name (act 280). If there are no additional characters, then module 164
14 outputs the decasified name as the *L* string and the corresponding case information
15 as the *I* string (act 282). However, if there are additional characters, then module
16 164 selects the next character in the input name (act 284) and returns to check
17 whether each of the leading eight bits of the selected character is zero (act 276).

18 Returning to act 276, if the leading eight bits of the selected character are
19 all zeroes, then module 164 checks whether the selected character is not a lower-
20 case character (act 286). It should be noted that the phrase "not a lower-case
21 character" is not equivalent to the phrase "an upper-case character". For example,
22 the asterisk symbol (*) has no case distinction and thus would not be a lower-case
23 character, but also would not be an upper-case character.

24 If the selected character is not a lower-case character, then the selected
25 character is appended to the *L* string and a zero bit is appended to the *I* string (act

288), and the process continues to act 280 to check whether any additional characters are in the input name. However, if the selected character is a lower-case character, then the upper-case form of the selected character is appended to the *L* string and a one bit is appended to the *I* string (act 290), and the process continues to act 280 to check whether any additional characters are in the input name.

By way of example, following the process of Fig. 7, if the input name were "Memo.doc", then module 164 would output as the decasified name (*L* string) "MEMO.DOC", and as the case information (*I* string) "01110111".

In the illustrated example of Fig. 7, a zero bit in case information string *I* is used to signify that a character is upper-case or has no case distinction, and the decasified name (*L* string) is all non-lower case characters. Various alterations can be made to this process, such as having upper-case characters signified by a one bit, having characters with no case distinction signified by a one bit, having the decasified name all non-upper case characters, etc.

Once decasified, the decasified name is encoded by encoding module (encoder) 166. Fig. 8 is a flowchart illustrating an exemplary process for encoding the decasified name in accordance with certain embodiments of the invention. Encoding module 166 initially receives the input decasified name (act 302). Encoding module 166 may receive the actual name, or alternatively an indication of where the name can be found (e.g., a pointer to the name), from module 160 or alternatively directly from decasifying module 164 (or mapping module 162 if the system is case-sensitive).

The characters of the input name are then reversed (act 304) and the number of trailing underscores in the reversed name counted and removed (act

306). An encoded bit stream *B* (which will be the output encoded name at the end of the encoding process) is then initialized with a string of one bits equal to the number of underscores removed in act 306, followed by a zero (act 308). If there are no trailing underscores, then the encoded bit stream *B* is initialized with a single zero bit. Alternatively, rather than relying on underscores, different characters may be used. In one implementation, the character that is Huffman coded to zero is the character that is counted and removed from the end of the reversed name (Huffman coding is discussed in more detail below).

Module 166 then checks whether the leading eight bits of the first character in the reversed name are all zero (act 310). If the leading eight bits are not all zero, then the sixteen bits of the first character are appended to the encoded bit stream *B* (act 312). However, if the leading eight bits of the first character are all zero, then the character (based on the last eight bits) is encoded using a first coding table (act 314), and then eight zero bits followed by the bits from the coding table are appended to the encoded bit stream *B* (act 316).

The coding table in act 314 can be generated in any of a wide variety of conventional manners using any of a wide variety of prefix coding schemes. In one implementation, conventional Huffman coding is used, although other coding schemes could alternatively be used. In the illustrated example, the coding table in act 314 codes an underscore to all zeroes and has codes for all valid non-lower case characters with Unicode values less than 256 and greater than 31, except for the following: space, period, quotation mark, asterisk, forward slash, colon, greater than symbol, less than symbol, question mark, back slash, or vertical line. If module 166 is given an input name with a character that is the first character in

act 308 and is not in the coding table, then the coding process fails, as does the modification process of Fig. 4.

Regardless of the bits added to encoded bit stream *B* in either act 312 or 316, module 166 proceeds to check whether there is an additional character in the reversed name (act 318). If there are no more characters in the reversed name, then module 166 removes any trailing zero bits and the one bit preceding the trailing zero bits from the encoded bit string *B* (act 320). Module 166 then outputs the encoded bit string *B* as the encoded name (act 322).

Returning to act 318, if there are additional characters in the reversed name, then the next character in the reversed name is selected (act 324) and a check made as to whether the leading eight bits of the selected character are all zero (act 326). If the leading eight bits are not all zero, then the sixteen bits of the selected character are appended to the encoded bit stream *B* (act 328). However, if the leading eight bits of the selected character are all zero, then the character (based on the last eight bits) is encoded using another coding table (act 330), and then eight zero bits followed by the bits from the coding table are appended to the encoded bit stream *B* (act 332).

Analogous to the coding table discussed in act 314 above, the coding table used in act 330 can be generated in any of a wide variety of conventional manners using any of a wide variety of prefix coding schemes. In one implementation, conventional Huffman coding is used, although other coding schemes could alternatively be used. In the illustrated example, the coding table in act 330 codes an underscore to all zeroes and has codes for all valid non-lower case characters with Unicode values less than 256 and greater than 31, except for the following: quotation mark, asterisk, forward slash, colon, greater than symbol, less than

1 symbol, question mark, back slash, or vertical line. These codes may be the same
2 as those in the table of act 314, or alternatively different. If module 166 is given
3 an input name with a character that is the selected character in act 324 and is not in
4 the coding table, then the coding process fails, as does the modification process of
5 Fig. 4.

6 Regardless of the bits added to encoded bit stream B in either act 328 or
7 332, module 166 proceeds to check whether there is an additional character in the
8 reversed name (act 318), and proceeds accordingly. In the illustrated example of
9 Fig. 8, each additional character is encoded using the other coding table in act 330.
10 Alternatively, the same coding table as was used in act 314 may be used in act
11 330. In another alternative, multiple additional coding tables are used in act 330
12 for encoding additional characters. For example, there may be a different coding
13 table for each additional character or groups of characters.

14 Once encoded, the encoded name (encoded bit string B of Fig. 8) is
15 encrypted by encryption module (encryptor) 168. Encryption module 168 initially
16 receives the encoded name (or an indication of where the name can be found) from
17 module 160 or alternatively directly from encoding module 166 (act 302).
18 Encryption module 168 can use any of a wide variety of encryption schemes to
19 encrypt the encoded name. The encryption scheme used, however, should be
20 surjective so that any arbitrary string of bits will decrypt to a valid encoded name,
21 and it should be deterministic so that identity of names can be determined by
22 examining only their encrypted forms. In the illustrated example, a conventional
23 block cipher is used, although other types of ciphers could alternatively be used
24 (e.g., a stream cipher).
25

1 Given that the encoded bit string B can be of varying length, additional bits
2 are added to the bit string B as necessary to bring the bit string up to a length that
3 is a multiple of the block size for the block cipher. In one implementation, this
4 addition (also referred to as padding) to bit string B is accomplished by prepending
5 to the bit string B (adding to the beginning of the bit string B) a one bit preceded
6 by as many zero bits as necessary to bring the length of the bit string up to a
7 multiple of the cipher block size. Encryption module 168 then encrypts the
8 padded bit string B on a per-block (x) basis according to the following three rules:

9 (1) if the block x of bit string B has the value of zero, then the encrypted
10 block has a value of zero;

11 (2) otherwise, if the result of encrypting block x with a block cipher
12 encryption E is zero, then the encrypted block has a value equal to the encryption
13 E of the value zero; and

14 (3) if neither rule (1) nor (2) applies, then the encrypted block has a value
15 equal to the encryption E of the block.

16 Any of a wide variety of conventional block ciphers can be used as block
17 cipher encryption E . In one implementation, block cipher encryption E is the
18 well-known RC2 block cipher. Other conventional block ciphers can alternatively
19 be used, such as DES, RC5, AES/Rijndael, and so forth.

20 It should also be noted that encryption module 168, using these rules, will
21 not result in an encrypted name with a first block equaling zero. This is true even
22 if the encryption uses a feedback mechanism such as cipher block chaining (CBC)
23 to increase the security of the encrypted data. In CBC mode, the plaintext of each
24 block is exclusively ORed (XORed) with the ciphertext of the previous block

1 before it is encrypted, thereby making each ciphertext block dependent not only on
2 the corresponding plaintext block but also on all previous plaintext blocks.

3 The resulting encrypted blocks are the encrypted name to be communicated
4 to another device (act 190 of Fig. 4).

5 In addition to encrypting the name, the case information *I* is also encrypted.
6 The case information *I* is encrypted using any of a wide variety of conventional
7 block ciphers. The block cipher may be the same as used for encrypting the
8 encoded name, or alternatively a different block cipher. Additionally, the case
9 information may be prepended with a one bit preceded by as many zero bits as
10 necessary to bring the length of the case information up to a multiple of the size of
11 the block cipher, analogous to the padding of the encoded bit string *B* discussed
12 above. Alternatively, additional random padding may be added to the end of the
13 case information *I* to bring it up to a multiple (not necessarily the next closest
14 multiple) of the block cipher size. Any such additional information is ignored
15 when recasifying the decoded name as discussed in more detail below.

16 Once generated, the encrypted directory entry is communicated to one or
17 more other computing devices 150. As discussed above, this is typically a
18 different device 150 than the device that generated the encrypted directory entry,
19 although it could be the same device. Regardless of the source of the encrypted
20 directory entry, the distributed system interface 200 at the receiving device 150
21 communicates the directory entry to server component 152.

22 Server component 152 includes a distributed system control module 350, a
23 syntax verifier 352, and a duplication identifier 354. Distributed system control
24 module 350 receives the encrypted directory entry and communicates the entry to
25 syntax verifier 352 to verify that the entry is syntactically legal, and duplication

1 identifier 354 to verify that the entry is not a duplicate of another entry in the
2 distributed file system. Distributed system control module 350 maintains a record
3 356 of the encrypted directory entries that are stored at computing device 150 (or
4 alternatively that are stored in all of the serverless distributed file system). In
5 addition to having record 356 available in memory 156, if memory 156 is a
6 volatile memory then record 356 is also stored in a nonvolatile memory (such as
7 mass storage device 158) in order to preserve record 356 if computing device 150
8 is powered-down. Additionally, distributed system control module 350
9 communicates with mass storage device 158 to store and retrieve encrypted files
10 corresponding to encrypted directory entries in record 356.

11 Syntax verifier 352 verifies that the syntax of a received directory entry is
12 valid. In the illustrated example, syntax verifier 352 checks whether the first
13 block of the encrypted directory entry is zero. If the first block is zero, then the
14 directory entry is not syntactically legal; otherwise, the directory entry is
15 syntactically legal.

16 Duplication identifier 354 compares the received encrypted directory entry
17 to the other encrypted directory entries stored in record 356. If the received
18 encrypted entry matches any of the encrypted entries stored in record 356 in the
19 same folder as the received encrypted entry will be placed in, then the received
20 entry is a duplicate; otherwise the received entry is not a duplicate. Note that the
21 serverless distributed file system prevents multiple files or folders within the same
22 directory or folder from having the same name, but does not prevent different files
23 in different folders from having the same name. For example, two files in the
24 same directory could not be named "memo.doc", but two files in two different
25

directories could be named "memo.doc". Record 356 maintains an indication, for each entry, of which folder or directory the entry is in.

Fig. 9 is a flowchart illustrating an exemplary process for verifying a received directory entry in accordance with certain embodiments of the invention. Initially, the encrypted name is received by distributed system control module 350 (act 370). Syntax verifier 352 then checks whether the first block of the encrypted name is zero (act 372). If the first block is zero, then the encrypted name is invalid (act 374). However, if the first block is not zero, then duplication identifier 354 checks whether the encrypted name is identical to any other encrypted name in the directory (act 376). If the encrypted name is identical to another encrypted name in the directory, then the encrypted name is invalid (act 374); otherwise, the encrypted name is valid (act 380).

As mentioned above, the exclusive encryption described herein maintains the following three properties:

- no two encrypted entries in a directory will decrypt to the same name
- all encrypted entries in the directory will decrypt to syntactically legal names
- the service that maintains the directory does not have access to the plaintext names of the entries

Distributed system control module 350 ensures that no two entries in a directory can be duplicates (based on the comparison performed by duplication identifier 354). Encryption of the directory entry using the block or other cipher ensures that the service that maintains the directory does not have access to the plaintext names of the entries.

1 The encryption process does not result in an encrypted name having a first
2 block that equals zero. Thus, any encrypted name having a first blocker equal to
3 zero is an illegal entry and is detected by syntax verifier 352. Additionally, in one
4 implementation a directory entry must satisfy the following criteria in order to be
5 syntactically legal:

- 6 • the directory entry name cannot be null (the encoding process
7 discussed above with reference to Fig. 8 admits no null name)
- 8 • the directory entry name cannot include any character with an ASCII
9 or Unicode value less than 32 (the encoding process discussed above
10 with reference to Fig. 8 admits no such characters)
- 11 • the directory entry name may not include any of the following
12 characters: " * / : < > ? \ | (the encoding process discussed above
13 with reference to Fig. 8 can be configured to admit no such
14 characters)
- 15 • the final character of a directory entry name may not be a space or a
16 period (the encoding process discussed above with reference to Fig.
17 8 can be configured to admit no such final characters)
- 18 • the following names are illegal (where *n* is any digit): CONIN\$,
19 CONOUT\$, CON, PRN, AUX, NUL, LPT n COM n (the mapping
20 process discussed above with reference to Fig. 6 can be configured
21 to admit no such names)
- 22 • for purposes of determining duplication, character case is ignored
23 (the decasifying process discussed above with reference to Fig. 7
24 decasifies the entry names)

Given that the directory entries are stored in an encrypted manner, when a computing device 150 retrieves the entries from another device in the serverless distributed file system, the entries need to be converted to plaintext to be useable by other components in the device. The general process for converting the received encrypted name and case information into a plaintext name is discussed above with reference to Fig. 5. One implementation of the process will be described in more detail below.

Initially, the encrypted name (bit string B discussed above) and the encrypted case information are decrypted by decryption module (decryptor) 170. The decryption process is the reverse of the encryption process discussed above. Initially, the prepended encoded name B is decrypted on a per-block (x) basis according to the following rules:

(1) if the block x of bit string B has the value of zero, then the decrypted block has a value of zero;

(2) otherwise, if the result of decrypting block x with a block cipher decryption E^{-1} is zero, then the decrypted block has a value equal to the decryption E^{-1} of the value zero; and

(3) if neither rule (1) nor (2) applies, then the decrypted block has a value equal to the decryption E^{-1} of the block.

The block cipher decryption E^{-1} is the decryption that corresponds to block cipher encryption E .

The result of this decryption process is the prepended encoded name B , from which decryption module 170 discards all leading bits up to and including the first one bit.

Decryption module 170 decrypts case information analogous to the decryption of prepended encoded name *B*. However, if any additional random padding was added to the end of the case information prior to encryption, this additional padding remains part of the decrypted case information and is simply ignored during the recasification process discussed below.

Once decrypted, decoding module (decoder) 172 decodes the encoded name *B*. Fig. 10 is a flowchart illustrating an exemplary process for decoding an encoded name in accordance with certain embodiments of the invention. Initially, decoding module 172 receives the encoded name (act 400), either from control module 160 or directly from decryption module 170. Decoding module 172 then initializes a name string *L* to null (act 402). The name string *L* will be the decoded name output by module 172 upon completion of the process of Fig. 10. The encoded name is also extended, if necessary, by appending a one bit followed by as many zero bits as needed to complete the process described in Fig. 10.

The number of leading one bits that precede the first zero bit in the encoded name *B* are then counted (act 404) and a pointer initialized to point to the first bit of the encoded name following the zero bit that follows the leading one bit (act 406). Decoding module 172 then checks whether the leading eight bits pointed to by the pointer are all zeroes (act 408). If the leading eight bits are not all zeroes, then module 172 appends the leading sixteen bits pointed to by the pointer to the name string *L* (act 410), updates the pointer to point to the first bit after the leading sixteen bits (act 412), and proceeds to check whether the leading eight bits pointed to by the pointer are all zeroes (act 414).

Returning to act 408, if the leading eight bits pointed to by the pointer are all zeroes, then module 172 decodes the next eight bits after the leading eight bits

1 using a coding table (act 416). The coding table used in act 416 is the same as the
 2 coding table used in act 314 of Fig. 8. Decoding module 172 then checks whether
 3 the decoded character is an underscore (act 418). If the decoded character is an
 4 underscore, then module 172 checks whether any more one bits remain in the
 5 encoded name *B* (act 420). If no more one bits remain, then decoding module 172
 6 appends a string of underscore characters to *L* equal in number to the number of
 7 leading one bits that preceded the first zero bit as counted in act 404 above (act
 8 422). The order of the characters in *L* is then reversed (act 424) and *L* is output as
 9 the decoded name (act 426).

10 Returning to acts 414 and 418, if the decoded character is not an underscore
 11 (act 418) or more one bits remain in the encoded name (act 420), then the decoded
 12 character is appended to *L* (act 428). Decoding module 172 then updates the
 13 pointer to point to the next bit after the decoded character (act 430), and then
 14 checks whether the leading eight bits pointed to by the pointer are all zeroes (act
 15 414). If the leading eight bits pointed to by the pointer are not all zeroes, then
 16 module 172 returns to act 410. However, if the leading eight bits pointed to by the
 17 pointer are all zeroes, then module 172 decodes the next eight bits after the leading
 18 eight bits using another coding table (act 432) and returns to act 418. The coding
 19 table used in act 432 is the same as the coding table used in act 330 of Fig. 8
 20 above.

21 Upon completion of the decoding process, a decoded and decasified name
 22 remains, along with corresponding case information. Recasifying module
 23 (recasifier) 174 uses the decoded name and the corresponding case information to
 24 recasify the name. Fig. 11 is a flowchart illustrating an exemplary recasifying
 25 procedure in accordance with certain embodiments of the invention. Initially,

recasifying module 174 receives an input name and corresponding case information (act 450). The input name and/or case information can be received from control module 160, or alternatively directly from decoding module 172.

Recasifying module 174 then initializes a decoded name string *M* to be a null bit string (act 452). Recasifying module 174 outputs string *M* as the recasified name upon completing the process of Fig. 11. Recasifying module 174 then selects the first character of the input name and the first bit of case information (act 454) and checks whether the selected case information bit is zero (act 456).

If the selected case information bit is zero, then the selected character is appended to string *M* (act 458). However, if the selected case information is not zero, then module 174 checks whether the leading eight bits of the selected character are all zeroes (act 460). If the leading eight bits of the selected character are not all zero, then the selected character is appended to string *M* (act 458); otherwise, module 174 checks whether the selected character has case distinction (act 462). If the selected character does not have case distinction then the selected character is appended to string *M* (act 458); otherwise, the lower-case form of the selected character is appended to string *M* (act 464).

After acts 458 and 464, module 174 checks whether any additional characters are in the input name (act 466). If there are no more characters in the input name then the value of string *M* is output as the recasified name (act 468). Any bits remaining in the case information are simply ignored. However, if there are additional characters in the input name, then module 174 selects the next character in the input name (act 470) and checks whether there is case information corresponding to the selected character (act 472). If there is case information

1 corresponding to the selected character, then the corresponding case information is
2 selected (act 474) and module 174 returns to act 456; otherwise, the value zero is
3 used as the selected case information corresponding to the selected character (act
4 476) and module 174 returns to act 456.

5 After recasifying module 174 recasifies the decode name, demapping
6 module (demapper) 176 demaps the recasified name. Fig. 12 is a flowchart
7 illustrating an exemplary process for demapping a name in accordance with
8 certain embodiments of the invention. Initially, an input name is received by
9 demapping module 176 (act 500). The input name can be received from control
10 module 160, or alternatively directly from recasifying module 174 (or alternatively
11 directly from decoding module 172 if the serverless distributed file system is case
12 sensitive).

13 Demapping module 176 then checks whether the input name is equal to an
14 illegal name followed by zero or more underscores (act 502). If the input name is
15 equal to an illegal name followed by zero or more underscores, then the demapped
16 name is formed by appending one underscore to the input name (act 504).
17 However, if the input name is not equal to an illegal name followed by zero or
18 more underscores, then the input name is used as the demapped name (act 506).
19 Thus, if the input name is not equal to an illegal name followed by zero or more
20 underscores, then the input name demaps to itself.

21 The system and process described above with reference to Figs. 3 – 12
22 discusses generation and storage of modified names that reveal very little
23 information about the original plaintext name (the approximate length of the
24 original plaintext name and possibly some case information). In an alternate
25 embodiment, the system and process described above are altered so that less

information is revealed about the original plaintext name. These alterations are to the encoding/decoding, encryption/decryption, and validation processes as follows.

The decoding and encoding processes (discussed above with reference to Figs. 8 and 10) and corresponding modules 166 and 172 of Fig. 3 are altered so that in act 320 of Fig. 8, only the trailing zero bits (if any) are removed (the one bit preceding the trailing zero bits is not removed). Then, during the decoding process of Fig. 10, the encoded name is extended by appending as many zero bits as necessary to complete the process described in Fig. 10.

Encryption module 168 of Fig. 3 is altered to the following encryption scheme. A fixed value W is established for the count of blocks in an encrypted name. This imposes some restrictions on the plaintext name (a plaintext name that would result in ciphertext having more than W blocks cannot be encrypted using this scheme). The encryption scheme can use any conventional block cipher encryption E (as discussed above) having a block size of Q bits. Encryption module 168 encrypts the encoded bit string according to the following three rules:

(1) if the encoded bit string B equals $W \times Q$ zero bits followed by a single one bit, then remove the trailing one bit from string B and encrypt the remaining bits of string B with block cipher E ;

(2) if rule (1) does not apply and if the length of the encoded bit string B is greater than $W \times Q$, then the procedure is aborted (the bit string is too long to be encrypted); and

(3) if neither rule (1) nor (2) applies, then append as many zero bits as necessary to the end of encoded bit string B so that bit string B has a length of $W \times Q$ bits, and encrypt the bit string B with block cipher E .

Encryption module 168 is also altered to encrypt the case information I as follows. The case information I is extended to include as many bits as necessary to bring the total length of the case information string I up to a fixed value V . The value V can vary, and in one implementation is equal to the largest number of characters that can be in a valid name (optionally plus additional space for additional random padding). The case information I can be extended using any bit values (e.g., random values). The case information I is then encrypted with block cipher E (or alternatively a different block cipher).

Decryption module 170 is similarly modified so that the decryption is performed as follows. Decryption module 170 sets the encoded name B equal to the decryption E^{-1} of the encrypted name. If the encoded name thus decrypted is all zero bits, then a one bit is appended to encoded name B . The encrypted case information is similarly decrypted using the decryption corresponding to the block cipher encryption used to encrypt the case information I .

The validation process (discussed above with reference to Fig. 9 and modules 350 and 352 of Fig. 3) is similarly altered to eliminate the checking of whether the first block of an encrypted name equals zero (block 372 of Fig. 9). The encryption process described here with a fixed block count W can result in valid encrypted names with a first block having a value of zero, so this checking is removed from the validation process.

The system and process described above with reference to Figs. 3 – 12 allows multiple properties to be maintained (no two entries in a directory can be encryptions of identical names, all entries in the directory are encryptions of syntactically legal names, and the service that maintains the directory does not have access to the plaintext names of the entries). Alternative systems and

processes may also be used that involve fewer acts and/or components, but do not necessarily maintain all of the properties discussed above. One such alternative is to simply encode a plaintext name using a fixed-bit-width encoding scheme, and then encrypt the encoded name using an encryption key. This alternative, however, can have problems when used with some encryption ciphers if the number of characters that are to be encoded is not a power of two (so there is no fixed bit width that can exactly encode all legal characters). Another such alternative also uses a fixed-bit-width encoding scheme, but with some characters encoded multiple times. This alternative, however, under circumstances cannot ensure that two directory entries are not the same.

Another such alternative is to encode each plaintext character using Huffman coding (or other prefix coding) and then encrypt the encoded name. However, due to variances in lengths of encoded names, it is not always possible to determine whether an encrypted name represents a legally complete encoded name. Another alternative that solves part of this problem is to discard any final partial character during decryption or to treat all missing bits as zeroes for the final character when running out of characters during decryption. However, both of these can also result in multiple directory entries that are the same.

Fig. 13 illustrates an example of a suitable operating environment in which the invention may be implemented. The illustrated operating environment is only one example of a suitable operating environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Other well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems,

0976496-01101
T04T0294946

1 microprocessor-based systems, programmable consumer electronics (e.g., digital
2 video recorders), gaming consoles, cellular telephones, network PCs,
3 minicomputers, mainframe computers, distributed computing environments that
4 include any of the above systems or devices, and the like.

5 Fig. 13 shows a general example of a computer 542 that can be used in
6 accordance with the invention. Computer 542 is shown as an example of a
7 computer in which various embodiments of the invention can be practiced, and
8 can be, for example, a computing device 150 of Fig. 3, a computing device 120 –
9 126 of Fig. 2, a computing device 102 or 104 of Fig. 1, etc. Computer 542
10 includes one or more processors or processing units 544, a system memory 546,
11 and a bus 548 that couples various system components including the system
12 memory 546 to processors 544.

13 The bus 548 represents one or more of any of several types of bus
14 structures, including a memory bus or memory controller, a peripheral bus, an
15 accelerated graphics port, and a processor or local bus using any of a variety of
16 bus architectures. The system memory 546 includes read only memory (ROM)
17 550 and random access memory (RAM) 552. A basic input/output system (BIOS)
18 554, containing the basic routines that help to transfer information between
19 elements within computer 542, such as during start-up, is stored in ROM 550.
20 Computer 542 further includes a hard disk drive 556 for reading from and writing
21 to a hard disk, not shown, connected to bus 548 via a hard disk drive interface 557
22 (e.g., a SCSI, ATA, or other type of interface); a magnetic disk drive 558 for
23 reading from and writing to a removable magnetic disk 560, connected to bus 548
24 via a magnetic disk drive interface 561; and an optical disk drive 562 for reading
25 from and/or writing to a removable optical disk 564 such as a CD ROM, DVD, or

1 other optical media, connected to bus 548 via an optical drive interface 565. The
2 drives and their associated computer-readable media provide nonvolatile storage
3 of computer readable instructions, data structures, program modules and other data
4 for computer 542. Although the exemplary environment described herein employs
5 a hard disk, a removable magnetic disk 560, and a removable optical disk 564, it
6 will be appreciated by those skilled in the art that other types of computer readable
7 media which can store data that are accessible by a computer, such as magnetic
8 cassettes, flash memory cards, random access memories (RAMs), read only
9 memories (ROM), and the like, may also be used in the exemplary operating
10 environment.

11 A number of program modules may be stored on the hard disk, magnetic
12 disk 560, optical disk 564, ROM 550, or RAM 552, including an operating system
13 570, one or more application programs 572, other program modules 574, and
14 program data 576. A user may enter commands and information into computer
15 542 through input devices such as keyboard 578 and pointing device 580. Other
16 input devices (not shown) may include a microphone, joystick, game pad, satellite
17 dish, scanner, or the like. These and other input devices are connected to the
18 processing unit 544 through an interface 568 that is coupled to the system bus
19 (e.g., a serial port interface, a parallel port interface, a universal serial bus (USB)
20 interface, etc.). A monitor 584 or other type of display device is also connected to
21 the system bus 548 via an interface, such as a video adapter 586. In addition to the
22 monitor, personal computers typically include other peripheral output devices (not
23 shown) such as speakers and printers.

24 Computer 542 operates in a networked environment using logical
25 connections to one or more remote computers, such as a remote computer 588.

1 The remote computer 588 may be another personal computer, a server, a router, a
2 network PC, a peer device or other common network node, and typically includes
3 many or all of the elements described above relative to computer 542, although
4 only a memory storage device 590 has been illustrated in Fig. 13. The logical
5 connections depicted in Fig. 13 include a local area network (LAN) 592 and a
6 wide area network (WAN) 594. Such networking environments are commonplace
7 in offices, enterprise-wide computer networks, intranets, and the Internet. In
8 certain embodiments of the invention, computer 542 executes an Internet Web
9 browser program (which may optionally be integrated into the operating system
10 570) such as the "Internet Explorer" Web browser manufactured and distributed by
11 Microsoft Corporation of Redmond, Washington.

12 When used in a LAN networking environment, computer 542 is connected
13 to the local network 592 through a network interface or adapter 596. When used
14 in a WAN networking environment, computer 542 typically includes a modem 598
15 or other means for establishing communications over the wide area network 594,
16 such as the Internet. The modem 598, which may be internal or external, is
17 connected to the system bus 548 via a serial port interface 568. In a networked
18 environment, program modules depicted relative to the personal computer 542, or
19 portions thereof, may be stored in the remote memory storage device. It will be
20 appreciated that the network connections shown are exemplary and other means of
21 establishing a communications link between the computers may be used.

22 Computer 542 also includes a broadcast tuner 600. Broadcast tuner 600
23 receives broadcast signals either directly (e.g., analog or digital cable
24 transmissions fed directly into tuner 600) or via a reception device (e.g., via
25 antenna or satellite dish).

Computer 542 typically includes at least some form of computer readable media. Computer readable media can be any available media that can be accessed by computer 542. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules, or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory, or other memory technology; CD-ROM, digital versatile disks (DVD), or other optical storage; magnetic cassettes, magnetic tape, magnetic disk storage, or other magnetic storage devices; or any other media which can be used to store the desired information and which can be accessed by computer 542. Communication media typically embodies computer readable instructions, data structures, program modules, or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared, and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

The invention has been described in part in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines,

programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Typically the functionality of the program modules may be combined or distributed as desired in various embodiments.

For purposes of illustration, programs and other executable program components such as the operating system are illustrated herein as discrete blocks, although it is recognized that such programs and components reside at various times in different storage components of the computer, and are executed by the data processor(s) of the computer.

Alternatively, the invention may be implemented in hardware or a combination of hardware, software, and/or firmware. For example, one or more application specific integrated circuits (ASICs) could be designed or programmed to carry out the invention.

Conclusion

Although the description above uses language that is specific to structural features and/or methodological acts, it is to be understood that the invention defined in the appended claims is not limited to the specific features or acts described. Rather, the specific features and acts are disclosed as exemplary forms of implementing the invention.